

Reducing Order Enforcement Cost in Complex Query Plans

Technical Report

By

Ravindra N. Guravannavar

S. Sudarshan

Ajit A. Diwan

Ch. Sobhan Babu

Department of Computer Science and Engineering
Indian Institute of Technology Bombay
Mumbai

Abstract

Algorithms that exploit sort orders are widely used to implement joins, grouping, duplicate elimination and other set operations. Query optimizers traditionally deal with sort orders by using the notion of interesting orders. The number of interesting orders is unfortunately factorial in the number of participating attributes. Optimizer implementations use heuristics to prune the number of interesting orders, but the quality of the heuristics is unclear. Increasingly complex decision support queries and increasing use of covering indices, which provide multiple alternative sort orders for relations, motivate us to better address the problem of optimization with interesting orders.

We show that even a simplified version of optimization with sort orders is NP-hard and provide principled heuristics for choosing interesting orders. We have implemented the proposed techniques in a Volcano-style cost-based optimizer, and our performance study shows significant improvements in estimated cost. We also executed our plans on a widely used commercial database system, and on PostgreSQL, and found that actual execution times for our plans were significantly better than for plans generated by those systems in several cases.

1 Introduction

Decision support queries, extract-transform-load (ETL) operations, data cleansing and integration often use complex joins, aggregation, set operations and duplicate elimination. Sorting based query processing algorithms for these operations are well known. Sorting based algorithms are quite attractive when physical sort orders of one or more base relations fulfill the sort order requirements of operators either completely or partially. Further, secondary indices that cover a query¹ are being increasingly used in read-mostly environments. Query covering indices make it very efficient to obtain desired sort orders without accessing the data pages. These factors make it possible for sort based plans to significantly outperform hash based counterparts.

The notion of interesting orders [SMD⁺79] has allowed optimizers to consider plans that could be locally sub-optimal, but produce orders that are beneficial for other operators, and thus produce a better plan overall. However, the number of interesting orders for most operators is factorial in the number of attributes involved. This is not acceptable as queries in the aforementioned applications do contain large number of attributes in joins and set operations.

In this paper we consider the problem of optimization taking sort orders into consideration. We make the following technical contributions:

1. Often order requirements of operators are partially satisfied by inputs. For instance, consider a merge-join with join predicate ($r.c_1 = s.c_1$ and $r.c_2 = s.c_2$). A clustering index on $r.c_1$ (or on $r.c_2$ or $s.c_1$ or $s.c_2$) is helpful in getting the desired order efficiently; a secondary index that covers the query has the same effect.

We highlight (in Section 3) the need for exploiting partial sort orders and show how a minor modification to the standard replacement selection algorithm can avoid run generation I/O completely when input is known to have a partial sort order. Further, we extend a cost-based optimizer to take into account partial sort orders.

2. We consider operators with flexible order requirements and address the problem of choosing good interesting orders so that complete or partial sort orders already available from inputs can be exploited.
 - In Sections 4 we show that a special case of finding optimal sort orders is NP-hard and give a 2-approximation algorithm to choose interesting sort orders for a join tree.
 - In Section 5 we address a more general case of the problem. In many cases, the knowledge of indices and available physical operators in the system allows us to narrow down the search space to a small set of orders. We formalize this idea (in Section 5.1) through the notion of *favorable orders*, and propose a heuristic to efficiently enumerate a small set of promising sort orders. Unlike heuristics used in optimizer implementations, our approach takes into account issues such as (i) added choices of sort orders for base relations due to the use of query covering indices (ii) sort orders that partially match an order requirement (iii) requirement of same sort

¹*i.e.*, contain all attributes of the relation that are used in the query

order from multiple inputs (*e.g.*, merge based join, union) and (*iv*) common attributes between multiple joins, grouping and set operations.

In Section 5.2 we also show how to integrate our extensions into a cost-based optimizer.

3. We present experimental results (in Section 6) evaluating the benefits of the proposed techniques. We compare the plans generated by our optimizer with those of three widely used database systems and show significant benefits due to each of our optimizations.

2 Related Work

Both System R [SMD⁺79] and Volcano [GM93] optimizers consider plans that could be locally sub-optimal but provide a sort order of interest to other operators and thus yield a better plan overall. However, the papers assume operators have one or few *exact* sort orders of interest. This is not true of operators like merge-join, merge-union, grouping and duplicate elimination that have a factorial number of interesting orders. Heuristics such as the PostgreSQL heuristic we shall see, are commonly used by optimizers. Details of the heuristics are publicly available only for PostgreSQL. Further, System R and Volcano optimizers consider only those sort orders as useful that completely meet an order requirement. Plans that partially satisfy a sort order requirement are not handled. In this report we address these two issues.

The seminal work by Simmen et.al. [SSM96] describes techniques to infer orders from functional dependencies and predicates applied and thereby avoids redundant sort enforcers in the plan. The paper briefly mentions the problem of non-exact sort order requirements and mentions an approach of propagating an order specification that allows any permutation on the attributes involved. Though such an approach is possible for single input operators like group-by, it cannot be used for operators such as merge-join and merge-union for which the order guaranteed by both inputs must match. Moreover, the paper does not make it clear how the flexible order requirements are combined at other joins and group-by operators. Simmen et.al. [SSM96] mention that the approach of carrying a flexible order specification also increases the complexity of the code significantly. Our techniques do not use flexible order specifications and hence can be incorporated into an existing optimizer with minimal changes. Further, our techniques work uniformly across all types of operators that have a flexible order requirement. Work on inferring orders and groupings [SSM96] [WC03] [NM04a] [NM04b] is independent and complementary to our work.

3 Exploiting Partial Sort Orders

Often, sort order requirements of operators are partially satisfied by indices or other operators in the input subexpressions. A prior knowledge of partial sort orders available from inputs allows us to efficiently produce the required (complete) sort order more efficiently. When operators have flexible order requirements, it is thus important to choose a sort order that makes maximum use

of partial sort orders already available. We motivate the problem with an example. Consider the query shown in Example 1. Such queries frequently arise in consolidating data from multiple sources. The join predicate between the two *catalog* tables involves four attributes and two of these attributes are also involved in another join with the *rating* table. Further, the order-by clause asks for sorting on a large number of columns including the columns involved in the join predicate.

Example 1 *A query with complex join condition*

```
SELECT c1.make, c1.year, c1.city, c1.color, c1.sellreason, c2.breakdowns, r.rating
FROM catalog1 c1, catalog2 c2, rating r
WHERE c1.city=c2.city AND c1.make=c2.make AND c1.year=c2.year AND c1.color=c2.color AND
      c1.make=r.make and c1.year=r.year
ORDER BY c1.make, c1.year, c1.color, c1.city, c1.sellreason, c2.breakdowns, r.rating;
```

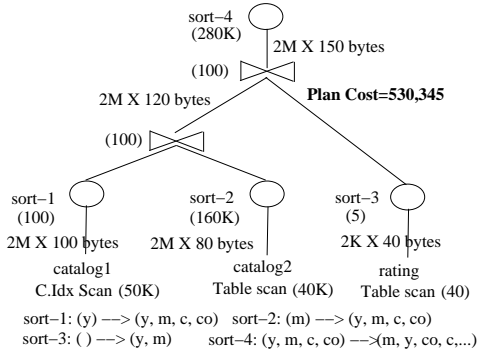


Figure 1: A naïve plan

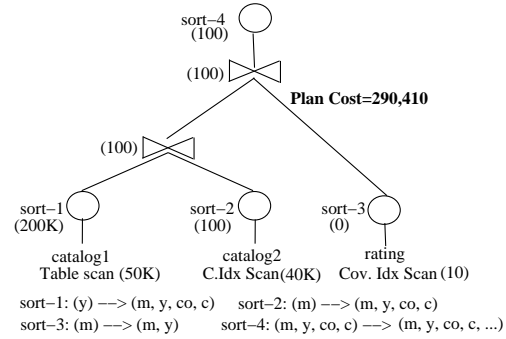


Figure 2: Optimal merge-join plan

The two catalog tables contain 2 million records each and have average tuple sizes of 100 and 80. We assume a disk block size of 4K bytes and 10000 blocks (40 MB) of main memory for sorting. The table *catalog1* is clustered on *year* and the table *catalog2* is clustered on *make*. The *rating* table has a secondary index on the *make* column with the *year* and *rating* columns included in the leaf pages (a covering index). Figure 1 and 2 show two different plans for the example query. Numbers in the parentheses indicate estimated cost of the operators in number of I/Os (CPU cost is appropriately translated into I/O cost units). Edges are marked with the number of tuples expected to flow on that edge and their average size. For brevity, the input and output orders for the sort enforcers are shown using only the starting letters of the column names. Though both plans use the same join order and employ sort-merge joins, the second plan is expected to perform significantly better than the first.

3.1 Changes to External Sort

External sorting algorithms have been studied extensively but in isolation. The standard replacement selection [Knu73] for run formation well adapts with the extent to which input is presorted. In the extreme case, when the input is fully sorted, it generates a single run on the disk and avoids

merging altogether. Larson [Lar03] revisits run formation in the context of query processing and extends the standard replacement selection to handle variable length keys and to improve locality of reference (reduced cache misses). Estivill-Castro and Wood [ECW92] provide a survey of adaptive sorting algorithms. The technique we propose in this section to exploit partial sort orders is a specific optimization in the context of multi-key external sorting. We observe that, by exploiting prior knowledge of partial sort order of input, it is possible to eliminate disk I/O altogether and have a completely pipelined execution of the sort operator.

We use the following notations: We use o, o_1, o_2 etc. to refer to sort orders. Each sort order o is a sequence of attributes/columns (a_1, a_2, \dots, a_n) . We ignore the sort direction (ascending/descending) as our techniques are applicable independent of the sort direction.

- ϵ : Empty (no) sort order
- $attrs(o)$: The set of attributes in sort order o
- $|o|$: Number of attributes in the sort order o
- $o_1 \leq o_2$: Order o_2 subsumes order o_1 (o_1 is a prefix of o_2)
- $o_1 < o_2$: Order o_1 is a strict prefix of o_2

Consider a case where the sort order to produce is (col_1, col_2) and the input already has the order (col_1) . Standard replacement-selection writes a single large run to the disk and reads it back again; this breaks the pipeline and incurs substantial I/O for large inputs. It is not difficult to see how the standard replacement-selection can be modified to exploit the partial sort orders. Let $o = (a_1, a_2, \dots, a_n)$ be the desired sort order and $o' = (a_1, a_2, \dots, a_k)$, $k < n$ be the partial sort order known to hold on the input. At any point during sorting we need to retain only those tuples that have the same value for attributes a_1, a_2, \dots, a_k . When a tuple with a new value for these set of attributes is read, all the tuples in the heap (or on disk if there are large number of tuples matching a given value of a_1, a_2, \dots, a_k) can be sent to the next operator in sorted order. Thus in most cases, partial sort orders allow a completely pipelined execution of the sort. Exploiting partial sort orders in this way has several benefits:

1. Let $o = (a_1, a_2, \dots, a_n)$ be the desired sort order and $o' = (a_1, a_2, \dots, a_k)$, $k < n$ be the partial sort order known to already hold on the input. We call the set of tuples that have the same value for attributes (a_1, a_2, \dots, a_k) as a *partial sort segment*. If each *partial sort segment* fits in memory (which is quite often the case in practice), the entire sort operation can be completed without any disk I/O.
2. Exploiting partial sort orders allows us to output tuples early (as soon as a new segment starts). In a pipelined execution this can have large benefits. Moreover, producing tuples early has immense benefits for Top-K queries and situations where the user retrieves only some result tuples.
3. Since sorting of each *partial sort segment* is done independently, the number of comparisons are significantly reduced. Note that we empty the heap every time a new segment

starts and hence insertions into heap will be faster. In general, independently sorting k segments each of size n/k elements, has the complexity $O(n \log(n/k))$ as against $O(n \log(n))$ for sorting all n elements. Further, while sorting each *partial sort segment* comparisons need to be done on fewer attributes, (a_{k+1}, \dots, a_n) in the above case.

Our experiments (Section 6) confirm that the benefits of exploiting partial sort orders can be substantial and yet none of the systems we evaluated, though widely used, exploited the partial sort orders.

3.2 Optimizer Extensions for Partial Sort Orders

In this section we assume order requirements of operators are concrete and only focus on incorporating partial sort orders. We deal with flexible order requirements in subsequent sections.

We use the following notations:

- $o_1 \wedge o_2$: Longest common prefix between o_1 and o_2
- $o_1 + o_2$: Order obtained by concatenating o_1 and o_2
- $o_1 - o_2$: Order o' such that $o_2 + o' = o_1$ (defined only when $o_2 \leq o_1$)
- $coe(e, o_1, o_2)$: The cost of enforcing order o_2 on the result of expression e which already has order o_1
- $N(e)$: Expected size, in number of tuples, of the result of expression e
- $B(e)$: Expected size, in number of blocks, of the result of expression e
- $D(e, s)$: Number of distinct values for attribute(s) s of expression e ($= N(\Pi_s(e))$)
- $cpu_cost(e, o)$: CPU cost of sorting result of e to get order o
- M : Number of memory blocks available for sorting

The *Volcano* optimizer framework [GM93] assumes an algorithm (physical operator) either guarantees a required sort order fully or it does not. Further, a physical property enforcer (such as sort) only knows the property to be enforced and has no information about the properties that hold on its input. The optimizer's cost estimate for the enforcer thus depends only on the required output property (sort order). In order to remedy these deficiencies we extended the optimizer in the following way: Consider an optimization goal (e, o) , where e is the expression and o the required output sort order. If the physical operator being considered for the logical operator at the root of e guarantees a sort order $o' < o$, then the optimizer adds a partial sort enforcer *enf* to enforce o from o' . We use the following cost model to account for the benefits of partial sorting.

$$coe(e, \epsilon, o) = \begin{cases} cpu_cost(e, o) & \text{if } B(e) \leq M \\ B(e)(2\lceil \log_{M-1}(B(e)/M) \rceil + 1) & \text{otherwise} \end{cases}$$

If e is known to have the order o_1 , we estimate the cost of obtaining an order o_2 as follows:

$coe(e, o_1, o_2) = D(e, attrs(o_s)) * coe(e', \epsilon, o_r)$, where $o_s = o_2 \wedge o_1$, $o_r = o_2 - o_s$ and $e' = \sigma_p(e)$, where p equates attributes in o_s to an arbitrary constant. Intuitively, we consider the cost of

sorting a single *partial sort segment* independently and multiply it by the number of segments. Note that we assume uniform distribution of values for $\text{attrs}(o_s)$. Therefore, we estimate $N(e') = N(e)/D(e, \text{attrs}(o_s))$ and $B(e') = B(e)/D(e, \text{attrs}(o_s))$. When the actual distribution of values is available, a more accurate cost model that does not rely on the uniform distribution assumption can be used.

4 Choosing Sort Orders for a Join Tree

Consider a join expression $e = e_1 \bowtie e_2$, where e_1, e_2 are input subexpressions and the join predicate is of the form: $(e_1.a_1 = e_2.a_1 \text{ and } e_1.a_2 = e_2.a_2 \dots \text{ and } e_1.a_n = e_2.a_n)$. Note that, *w.l.g.*, we use the same name for attributes being compared from either side and we call the set $\{a_1, a_2, \dots, a_n\}$ as the join attribute set. In this case, the merge join algorithm has potentially $n!$ interesting sort orders on inputs e_1 and e_2 ². The specific sort order chosen for the merge-join can have significant influence on the plan cost due to the following reasons: (i) Clustering and covering indices, indexed materialized views and other operators in the subexpressions e_1, e_2 can make one sort order much cheaper to produce than another. (ii) The merge-join produces the same order on its output as the one selected for its inputs. Hence, a sort order that helps another operator above the merge-join can help eliminate a sort or just have a partial sort. In this section we show that a special case of the the problem of choosing optimal sort orders for a tree of merge-joins is *NP-Hard* and provide a 2-approximate algorithm for the problem. In the next section, we describe our heuristics for a more general setting of the problem in which we make use of the proposed 2-approximate algorithm.

4.1 Finding Optimal is NP-Hard

Consider a join expression $e = R_1 \bowtie R_2 \bowtie R_3 \dots R_n$ and a specific join order tree for the expression. Consider a special case where all base relations and intermediate results are of the same size and no indices built on the base relations. Now, the problem of choosing optimal sort orders for each join requires us to choose permutations of join attributes such that we maximize the length of longest common prefixes of permutations chosen for adjacent nodes. Figure 3, shows an example and an optimal solution under the model where the benefit for an edge is the length of the longest common prefix between the permutations chosen for adjacent nodes and we maximize the total benefit. The join attribute set for each join node is shown in curly braces besides the node. Permutations chosen in the optimal solution are indicated with angle brackets and the number on each edge shows the benefit for that edge. Below we state the problem formally.

Problem 1 *Let T be a binary tree of order n , with vertex set $V(T)$ and edge set $E(T)$. Each node v_i ($i = 1, \dots, n$) is associated with an attribute set s_i . Find a sequence of permutations $p_1, p_2 \dots p_n$, where p_i is a permutation of set s_i , such that the benefit function \mathcal{F} is maximum.*

²We assume merge-join requires sorting on all attributes involved in the join predicate. We do not consider orders on subsets of join attributes since the additional cost incurred at merge-join matches the benefit of sorting a smaller subset of attributes.

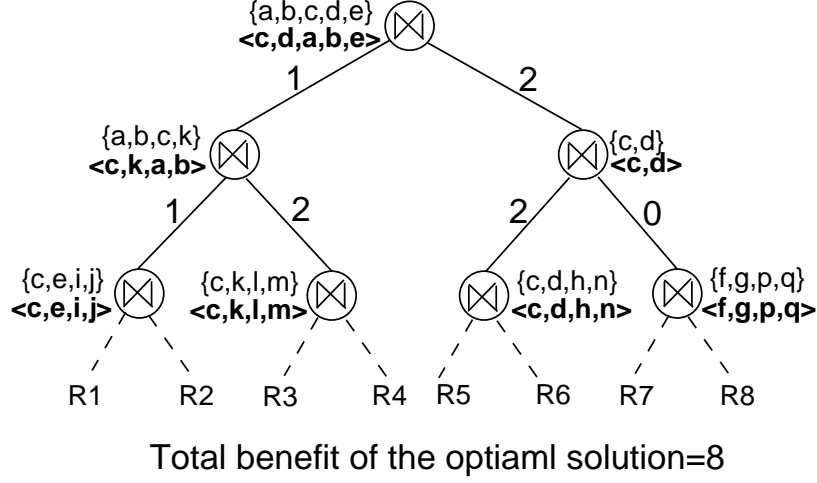


Figure 3: A special case of choosing globally optimal sort order

$$\mathcal{F} = \sum_{\forall v_i v_j \in E(T)} |p_i \wedge p_j|$$

Theorem 4.1 *Problem 1 is NP-hard.*

Proof: We give a reduction from the known *NP-Hard* problem *SUM-CUT* [DPS02].

Problem 2 (*SUM-CUT*) Given a graph $G = (V, E)$ with $m = |V|$ vertices, number the vertices of G as $1, \dots, m$ such that $\sum_{1 \leq i \leq m} c_i$ is minimized, where c_i is the number of vertices numbered $\leq i$ that are connected to at least one vertex numbered $> i$.

The *SUM-CUT* problem is equivalent to Problem 3 given below; to see the equivalence consider the complement of the graph.

Problem 3 Given a graph $G = (V, E)$ with $m = |V|$ vertices, number the vertices of G as $1, 2, \dots, m$ such that $\sum_{1 \leq i \leq m} q_i$ is maximized, where q_i is the number of vertices that are adjacent to *all* the vertices numbered 1 to i .

We now reduce Problem 3 to Problem 1. The construction is as follows:

Let the m vertices of the graph be labeled u_1, \dots, u_m . We construct the binary tree by choosing a vertex set $V(T)$ of size $2m$, $\{v_1, v_2, \dots, v_{2m}\}$. v_1, \dots, v_m are internal vertices and v_{m+1}, \dots, v_{2m} are leaf vertices. Let the edge set $E(T)$ be $\{v_i v_{i+1} : 1 \leq i < m\} \cup \{v_i v_{m+i} : 1 \leq i \leq m\}$. Attribute set s_i of internal vertex v_i is $V(G) \cup \mathcal{L}$, where \mathcal{L} is an arbitrarily large set disjoint from $V(G)$, for $1 \leq i \leq m$. Attribute set s_i of leaf vertex v_i is $\{w : w \in V(G) \text{ and } w u_{i-m} \in E(G)\}$, for $m < i \leq 2m$.

□

4.2 A 2-Approximate Algorithm

An efficient dynamic programming based algorithm to find the optimal solution (under the benefit model presented in the previous section) for Problem 1 exists when the tree is a *path*. Note that left-deep and right-deep join plans result in paths. We first present this algorithm and make use of it in finding a 2-approximation for binary trees.

Consider a path v_1, v_2, \dots, v_n , where each vertex v_i is associated with an attribute set s_i . The optimal solution for any segment (i, j) of the path, $\text{OPT}(i, j) = \max\{ \text{OPT}(i, k) + \text{OPT}(k + 1, j) + c(i, j) \}$ over all $i \leq k < j$, where $c(i, j)$ is the number of common attributes for the segment (i, j) .

The dynamic programming algorithm starts by finding the optimal solution for all segments of size 1, then of size 2 and so on. At any segment (i, j) , every possible way of splitting the segment into smaller segments is considered and the solution for the smaller segments obtained from the memo structure. The optimal solution for segment (i, j) is then memoized. Procedure *PathOrder* in Figure 4 gives the details.

For binary trees we propose an approximation with benefit at least half that of an optimal solution. We split the tree into two sets of paths, P_o and P_e . P_o has the paths formed by edges at odd levels and P_e has those formed by edges at even levels, Figure 5 shows an example. We use the procedure *PathOrder* to find optimal solutions for each of the two sets of paths. Let the optimal solutions for the two sets of paths be S_o and S_e and the corresponding benefits be $\text{ben}(S_o)$ and $\text{ben}(S_e)$. Let the set of edges included in P_o and P_e be denoted by E_o and E_e respectively. Consider an optimal solution S_T for the whole tree. In the optimal solution, let the sum of benefits of all edges in E_o be *odd-ben*(S_T) and that of edges in E_e be *even-ben*(S_T). Note that $\text{ben}(S_o) \geq \text{odd-ben}(S_T)$ and $\text{ben}(S_e) \geq \text{even-ben}(S_T)$. Since the total benefit of the optimal solution $\text{ben}(S_T) = \text{odd-ben}(S_T) + \text{even-ben}(S_T)$, we have $\text{ben}(S_o) + \text{ben}(S_e) \geq \text{ben}(S_T)$. Hence at least one of $\text{ben}(S_o)$ or $\text{ben}(S_e)$ is $\geq 1/2 \text{ben}(S_T)$. There may be vertices not included in the chosen solution, e.g., the even level split in Figure 5 does not include the root and leaf nodes. For these left over vertices arbitrary permutations can be chosen.

5 Optimization Exploiting Favorable Orders

The benefit model we presented in the previous section and the approximation algorithm do not take into account indices and size of relations or intermediate results. Moreover, we assumed that the join order tree is fixed. In this section we present a two phase approach for the more general problem. In phase-1, which occurs during plan generation, we exploit the information of available indices and properties of physical operators to efficiently compute a small set of promising sort orders to try. We formalize this idea through the notion of *favorable orders*. Phase-2, is a plan refinement step and occurs after the optimizer makes its choice of the best plan. In phase-2, the sort orders chosen by the optimizer are refined further to reap extra benefit from the attributes common to multiple joins. Phase-2 uses the 2-approximate algorithm of Section 4.2

Input:

$s[n]$: array of attribute sets

Output:

$p[n]$: array of permutations or orders

Data Structures:

$benefit[n][n]$, $split[n][n]$: arrays of integers

$commons[n][n]$: array of attribute sets

$apermute(s)$: Function that returns an arbitrary permutation of attribute set s

Procedure PathOrder

BEGIN

for $i=1$ to n

$benefit[i][i] = 0$; $commons[i][i] = s[i]$; $split[i][i] = -1$;

for $j=1$ to $n-1$

 for $i = 1$ to $n-j$

 Let k be the index such that $i \leq k < (i+j)$ and $benefit[i][k] + benefit[k+1][i+j]$ is maximum

$commons[i][i+j] = commons[i][k] \cap commons[k+1][i+j]$;

$benefit[i][i+j] = benefit[i][k] + benefit[k+1][i+j] + |commons[i][i+j]|$;

$split[i][i+j] = k$;

 Call MakePermutation($1, n$);

END PROC

Procedure MakePermutation(i, j)

BEGIN

 if ($i = j$)

$p[i] = \text{Append } apermute(commons[i][i]) \text{ to } p[i]$;

 return;

 for $k=i$ to j

$p[k] = \text{Append } apermute(commons[i][j]) \text{ to } p[k]$;

 For all $(i', j') \neq (i, j)$

$commons[i'][j'] = commons[i][j] - commons[i][j]$;

$m = split[i][j]$;

 MakePermutation(i, m);

 MakePermutation($m+1, j$);

END PROC

Figure 4: Optimal Orders for Path

5.1 Favorable Orders

Given an expression e , we expect some sort orders (on the result of e) to be producible at much lesser cost than other sort orders. Available indices, indexed materialized views, specific rewriting of the expression and choice of physical operators determine what sort orders are easy to produce. To account for such orders, we introduce the notion of *favorable* orders. We use the following notations:

- $cbp(e, o)$: Cost of the best plan for expression e with o being the required output order
- o_R : The clustering order of relation R
- $idx(R)$: Set of all indices over R
- $o(I)$: Order (key) of the index I




Figure 5: A 2-approximation for binary trees

- $\langle s \rangle$: An arbitrary order (permutation) on attribute set s
- $P(s)$: Set of all permutations of set s
- $o \wedge s$: Longest prefix of o such that each attribute in the prefix belongs to the attribute set s
- $schema(e)$: The set of attributes in the output of e

We first define the **benefit** of a sort order o w.r.t. an expression e as follows:

$$benefit(o, e) = cbp(e, \epsilon) + coe(e, \epsilon, o) - cbp(e, o)$$

Intuitively, a positive benefit implies the order can be obtained with lesser cost than a full sort of unordered result. For instance, the clustering order of a relation r will have a positive *benefit* for the expression $\sigma_p(r)$. Similarly, query covering secondary indices and indexed materialized views can yield orders with positive benefit. We call the set of all orders, on $schema(e)$, having a positive benefit w.r.t. e as the *favorable order set* of e and denoted it as $ford(e)$.

$$ford(e) = \{ o : benefit(o, e) > 0 \}$$

5.1.1 Minimal Favorable Orders

The number of favorable orders for an expression can be very large. For instance, every order having the clustering order as its prefix is a favorable order. A *minimal favorable order set* of e , denoted by $ford-min(e)$, is the minimum size subset of $ford(e)$ such that, for each order $o \in ford(e)$, at least one of the following is true:

1. o belongs to $ford-min(e)$
2. $\exists o' \in ford-min(e)$ such that $o' \leq o$ and $cbp(e, o') + coe(e, o', o) = cbp(e, o)$ Intuitively, if the cost of obtaining order o equals the cost of obtaining a partial sort order o' followed by an explicit sort to get o , we include only o' in the *ford-min*
3. $\exists o'' \in ford-min(e)$ such that $o \leq o''$ and $cbp(e, o'') = cbp(e, o)$ Intuitively, if an order o'' subsumes order o and has the same cost, we include only o'' in *ford-min*

Conditions 2 and 3 above, ensure that when a relation has an index on o , orders that are prefixes of o and orders that have o as their prefix are not included unnecessarily.

We define favorable orders of an expression *w.r.t.* a set of attributes s as: $\text{ford}(e, s) = \{o \wedge s : o \in \text{ford}(e)\}$. Intuitively, $\text{ford}(e, s)$ is the set of orders on s or a subset of s that can be obtained efficiently. Similarly, the *ford-min* of an expression *w.r.t.* a set of attributes s is defined as: $\text{ford-min}(e, s) = \{o \wedge s : o \in \text{ford-min}(e)\}$.

5.1.2 Heuristics for Favorable Orders

Note that the definition of favorable orders uses the cost of the best plan for the expression. However, we need to compute the favorable orders of an expression **before** the expression is optimized and without requiring to expand the logical or the physical plan space. Further, the size of the exact *ford-min* of an expression can be prohibitively large in the worst case. In this section, we describe a method of computing approximate *ford-min*, denoted as *afm*, for *SPJG* expressions. We compute the *afm* of an expression bottom-up. For any expression e , $\text{afm}(e)$ is computable after the *afm* is computed for all of e 's inputs.

1. $e = R$, where R is a base relation or materialized view. We include the clustering order of R and all secondary index orders such that the index covers the query.
 $\text{afm}(R) = \{o : o = o_R \text{ or } o = o(I), I \in \text{idx}(R) \text{ and } I \text{ covers the query}\}$
2. $e = \sigma_p(e_1)$, where e_1 is an arbitrary expression.
 $\text{afm}(e) = \{o : o \in \text{afm}(e_1)\}$
3. $e = \Pi_L(e_1)$, where e_1 is any expression. We include longest prefixes of input favorable orders such that the prefix has only the projected attributes.
 $\text{afm}(e) = \{o : \exists o' \in \text{afm}(e_1) \text{ and } o = o' \wedge L\}$
4. $e = e_1 \bowtie e_2$ with join attribute set $S = \{a_1, a_2, \dots, a_n\}$. Noting that nested loops joins propagate the sort order of one of the inputs (outer) and merge join propagates the sort order chosen for join attributes, we compute the *afm* as follows. First, we include all sort orders in the input *afms*. Next, we consider the longest prefix of each input favorable order having attributes only from the join attribute set and extend it to include an arbitrary permutation of the remaining join attributes.
 $\text{afm}(e_1 \bowtie e_2) = \mathcal{T} \cup \{o : o' \in \mathcal{T} \cup \{\epsilon\} \text{ and } o = o' \wedge S + \langle S\text{-attrs}(o' \wedge S) \rangle\}$, where $\mathcal{T} = \text{afm}(e_1) \cup \text{afm}(e_2)$

Note that, for the join attributes not involved in an input favorable order prefix (*i.e.*, $S\text{-attrs}(o' \wedge S)$), we take an arbitrary permutation. An exact *ford-min* would require us to include all permutations of such attributes. In the post-optimization phase, we refine the choice made here using the benefit model and algorithm of Section 4.2.

5. $e =_L \mathcal{G}_F(e_1)$
 $\text{afm}(e) = \{o : o' \in \text{afm}(e_1) \cup \{\epsilon\} \text{ and } o = o' \wedge L + \langle L\text{-attrs}(o' \wedge L) \rangle\}$
Intuitively, for each input favorable order we identify the longest prefix with attributes from the projected group-by columns and extend the prefix with an arbitrary permutation of the remaining attributes.

Computing *afms* requires a single pass of the query tree. At each node of the query tree the only significant operation performed is computation of the longest common prefix *lcp* of an order *w.r.t.* a set, *i.e.*, $o \wedge s$. The number of *lcp* operations performed at a node is at most the number of input favorable orders for that node. Let m be the average number of favorable orders for base relations. If the query involves a join of n relations the worst-case number of *lcp* operations performed at any node is $O(2^n m)$. Typically m is very small (≤ 2) and hence even the worst-case number of *lcp* operations is within acceptable limits. However, in most practical cases, the number of *lcp* operations performed is far below the worst-case.

5.2 Overall Optimizer Extensions

We make use of the approximate favorable orders during plan generation (phase-1) to choose a small set of promising *interesting orders* for sort-based operators. We describe our approach taking merge join as an example but the approach is applicable to other sort based operators. In phase-2, which is a post-optimization phase, we further refine the chosen sort orders.

5.2.1 Plan Generation (Phase-1)

Consider an optimization goal of expression $e = e_l \bowtie e_r$ and required output sort order o . When we consider merge-join as a candidate algorithm, we need to generate sub-goals for e_l and e_r with the required output sort order being some permutation of the join attributes.

Let S be the set of attributes involved in the join predicate. We consider only conjunctive and equality predicates. We compute the set $I(e, o)$ of interesting orders as follows:

1. Collect the favorable orders of inputs plus the required output order
 $\mathcal{T}(e, o) = afm(e_l, S) \cup afm(e_r, S) \cup o \wedge S$, where $afm(e, S) = \{o' \wedge S : o' \in afm(e)\}$
2. Remove redundant orders
 If $o_1, o_2 \in \mathcal{T}(e, o)$ and $o_1 \leq o_2$, remove o_1 from $\mathcal{T}(e, o)$
3. Compute the set $I(e, o)$ by extending each order in $\mathcal{T}(e, o)$ to the length of $|S|$; the order of extra attributes can be arbitrarily chosen
 $I(e, o) = \{o : o' \in \mathcal{T}(e, o) \text{ and } o = o' + \langle S - attrs(o') \rangle\}$

We then generate optimization sub-goals for e_l and e_r with each order $o' \in I(e, o)$ as the required output order and retain the cheapest combination.

A Note on Optimality: If the set $I(e, o)$ is computed using the exact *ford-mins* instead of *afms*, we claim that it must contain an optimal sort order (a sort order that produces the optimal merge join plan in terms of overall plan cost). Appendix A gives the detailed proof of this claim.

An Example: Consider Example 1 of Section 3. For brevity, we refer to the two catalog tables as *ct1* and *ct2*, the rating table as *rt* and the columns with their starting letters. The *afms* computed as described in Section 5.1.2 are as follows: $afm(ct1) = \{(y)\}$, $afm(ct2) = \{(m)\}$, $afm(rt) = \{(m)\}$, $afm(ct1 \bowtie ct2) = \{(y, co, c, m), (m, co, c, y)\}$, $afm((ct1 \bowtie ct2) \bowtie rt) = \{(y, m), (m, y)\}$

For $(ct1 \bowtie ct2) \bowtie rt$ we consider two interesting sort orders $\{(y, m), (m, y)\}$ and for $ct1 \bowtie ct2$ we consider the four orders $\{(y, co, c, m), (m, co, c, y), (y, m, co, c), (m, y, co, c)\}$. As a result the optimizer arrives at plan shown in Figure 2.

5.2.2 Plan Refinement (Phase-2)

During the plan refinement phase, for each merge-join node in the plan tree, we identify the set of *free attributes*, the attributes which were not part of any of the input favorable orders. Note that for these attributes we had chosen an arbitrary permutation while computing the *afm* (Section 5.1.2). We then make use of the 2-approximate algorithm for trees (Section 4.2) and rework the permutations chosen for the *free attributes*.

Formally, let p_i be the permutation chosen for the join node v_i . Let q_i be the order such that $q_i \in \text{afm}(v_i.\text{left-input}) \cup \text{afm}(v_i.\text{right-input})$ and $|p_i \wedge q_i|$ is the maximum. Intuitively, q_i is the input favorable order sharing the longest common prefix with p_i . Let $f_i = \text{attrs}(p_i - (p_i \wedge q_i))$; f_i is the set of *free attributes* for v_i .

We now construct a binary tree where each node n_i corresponding to join-node v_i is associated with the attribute set f_i . The orders for the nodes are chosen using the 2-approximate algorithm; the chosen order for free attributes is then appended to the order chosen during plan generation (i.e., $p_i \wedge q_i$) to get a complete order.

The reworking of the orders will be useful only if the adjacent nodes share the same prefix, i.e., $p_i \wedge q_i$ was the same for adjacent nodes. This condition however certainly holds when the inputs for joins have no favorable orders.

Figure 6 illustrates the post-optimization phase. Assume all relations involved ($R_1 \dots R_4$) are clustered on attribute a and no other favorable orders exist. i.e., $\text{afm}(R_i) = \{(a)\}$, for $i = 1$ to 4. The orders chosen by the plan generation phase are shown besides the join nodes with *free attributes* being in *italics*. The reworked orders after the post-optimization phase are shown underlined.




Figure 6: Post-Optimization Phase

6 Experimental Results

We performed experiments to evaluate the benefits our techniques. For comparison, we use PostgreSQL (version 8.1.3) and two widely used commercial database systems (we call them SYS1 and SYS2). All tests were run on an Intel P4 (HT) PC with 512 MB of RAM. We used TPC-H 1GB dataset and additional tables as specified in the individual test cases. For each table, a clustering index was built on the primary key. Additional secondary indices built are specified

along with the test cases. All relevant statistics were built and the optimization level for one of the systems, which supports multiple levels of optimization, was set to the highest.

6.1 Modified Replacement Selection

The first set of experiments evaluate the benefits of modified replacement selection (MRS) as compared to the standard replacement selection (SRS) when the input is known to be partially sorted.

External sort in PostgreSQL employs the standard replacement selection (SRS) algorithm suitably adapted for variable length records. We modified this implementation to exploit partial sort orders available on the input.

Experiment A1: The first experiment consists of a simple ORDER BY of the TPC-H *lineitem* table on two columns (*l_suppkey*, *l_partkey*).

Query 1 ORDER-BY on lineitem

```
SELECT l_suppkey, l_partkey FROM lineitem ORDER BY l_suppkey, l_partkey;
```

A secondary index on *l_suppkey* was available that covered the query (included the *l_partkey* column)³. On all three systems, the order by on (*l_suppkey*, *l_partkey*) took almost the same time as an order by on (*l_partkey*, *l_suppkey*) showing that the sort operator of these systems did not exploit partial sort orders effectively. We then compared the running times with our implementation that exploited partial sort order (*l_suppkey*) and the results are shown in Figure 7.

On SYS1 and SYS2 we simulated the partial sorting using a correlated rank query (as we did not have access to their source code). The subquery sorted the index entries matching a given *l_suppkey* on *l_partkey* and the subquery was invoked with all *suppkey* values so as to obtain the desired sort order of (*l_suppkey*, *l_partkey*).




Figure 7: Experiment 1




Figure 8: Rate of Output

³On systems not supporting indexes with included columns, we used a table with only the desired two columns, clustered on *l_suppkey*




Figure 9: Effect of Partial Sort Segment Size

By avoiding run generation I/O and making reduced comparisons, MRS performs 3-4 times better than SRS.

Experiment A2: The second experiment shows how MRS is superior in terms of its ability to produce records early and uniformly. Table R_3 having 3 columns (c_1, c_2, c_3) was populated with 10 million records and was clustered on (c_1). The query asked an order by on (c_1, c_2). Figure 8 shows the plot of number of tuples produced vs. time with cardinality of $c_1 = 10,000$.

MRS starts producing the tuples without any delay after the operator initialization where as SRS produces its first output tuple only after seeing all input tuples. By producing tuples early, MRS speeds up the pipeline significantly and also helps Top-K queries.

Experiment A3: The third experiment shows the effect of *partial sort segment size* on sorting. 8 tables $R_0 \dots R_7$, with identical schema of 3 columns (c_1, c_2, c_3) were each populated with 10 million records and average record size of 200 bytes. Each table was clustered on (c_1). Table R_i had 10^i tuples for each value of c_1 , resulting in a partial sort segment size of 200×10^i bytes. Thus R_0 had c_1 as unique and sort segment size of 200 bytes and R_7 had the same value of c_1 for all 10 million records leading to a sort segment size of 2GB. The query asked for an order by on (c_1, c_2). The running times with default and modified replacement selection on PostgreSQL are shown in Figure 9.

When the partial sort segment size is small enough to fit in memory (up to 10MB or 50K records), SRS produces a single sorted run on disk and does not involve merging of runs. The modified replacement selection (MRS) gets the benefit of avoiding I/O and reduced number of comparisons. When the partial sort segment size becomes too large to fit in memory, we see a sudden rise in the time taken by SRS. This is because replacement selection will have to deal with merging several runs. MRS however deals with merging smaller number of runs initially as each partial sort segment is sorted separately. As the partial sort segment size increases, the running time of MRS rises and becomes same as that of SRS at the extreme point where all records have the same value for c_1 .

Experiment A4: To see the influence of MRS on a query having other operators, we considered a query that asked for counting the number of lineitems for each supplier, part pair. Two indices, *lineitem(l_suppkey)* and *partsupp(ps_suppkey)*, each of which included other required columns supplied the required sort order partially.

Query 2 *Number of lineitems for each (supplier, part) pair*

```
SELECT ps_suppkey, ps_partkey, ps_availqty, count(l_partkey)
FROM   partsupp, lineitem
WHERE  ps_suppkey=l_suppkey AND ps_partkey=l_partkey
GROUP BY ps_suppkey, ps_partkey, ps_availqty
ORDER BY ps_suppkey, ps_partkey;
```

The query took 63 seconds to execute with SRS and 25 seconds with MRS, both on Postgres. The query plan used in both cases was the same - a merge join of the two relations on (*suppkey*, *partkey*) followed by an aggregate.

6.2 Choice of Interesting Orders

We extended our Volcano-style cost based optimizer, which we call PYRO, to consider partial sort orders and choose good interesting sort orders for merge joins and aggregation. We compare the plans produced by the extended implementation, which we call PYRO-O, with those of Postgres, SYS1 and SYS2.

Experiment B1: For this experiment we used Query 3 given below, which lists parts for which the outstanding order quantity is more than the stock available at the supplier.

Query 3 *Parts Running Out of Stock*

```
SELECT ps_suppkey, ps_partkey, ps_availqty, sum(l_quantity)
FROM   partsupp, lineitem
WHERE  ps_suppkey=l_suppkey AND ps_partkey=l_partkey AND l_linestatus='O'
GROUP BY ps_availqty, ps_partkey, ps_suppkey
HAVING sum(l_quantity) > ps_availqty ORDER BY ps_partkey;
```

Table *partsupp* had clustering index on its primary key (*ps_partkey*, *ps_suppkey*). Two secondary indices, one on *ps_suppkey* and the other on *l_suppkey* were also built on the *partsupp* and *lineitem* tables respectively. The two secondary indices covered all attributes needed for the query. The experiment shows the need for cost-based choice of interesting orders. The choice of interesting orders for the join and aggregate are not obvious in this case for the following reasons:

1. The explicit order by clause favors the choice of an order where *partkey* appears first.
2. The clustering index on *partsupp* favors the choice of (*partkey*, *suppkey*) as the interesting order.

3. The secondary indices favor the choice of $(suppkey, partkey)$ that can be obtained by using a low cost partial sort. Note that this option can be much cheaper due to the size of the *lineitem* relation.

Therefore, the optimizer must make a cost-based decision on the sort order to use. Figures 10 and 11 show the plans chosen by Postgres, PYRO-O, SYS1 and SYS2.

All plans except the hash-join plan of SYS1 and the plan produced by PYRO-O use an expensive full sort of 6 million *lineitem* index entries on $(partkey, suppkey)$. Further, Postgres uses a hash aggregate where a sort-based aggregate would have been much cheaper as the required sort order was available from the output of merge-join (note that the functional dependency $\{ps_partkey, ps_suppkey\} \rightarrow \{ps_availqty\}$ holds). On SYS1, it was possible to force the use of a merge-join instead of hash-join and the plan chosen is shown in Figure 11(b).

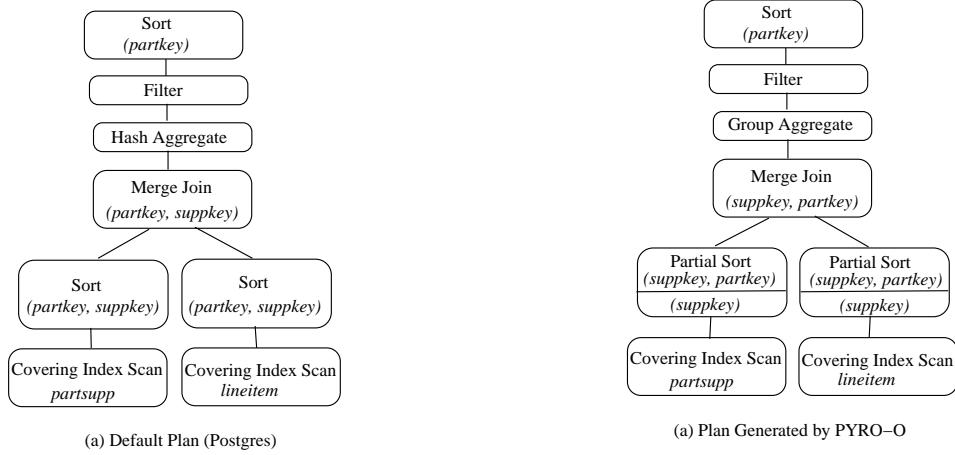



Figure 10: Plans for Query 3 (Postgres and PYRO-O)

We compared the actual running time of PYRO-O's plan with those of Postgres and SYS1 by forcing our plan on the respective systems. Figures 12 and 13 show the details. It was not possible for us to force our plan on SYS2 and make a fair comparison and hence we omit the same. The only surprising result was the default plan chosen by SYS1 performed slightly poorer than the forced merge-join plan on SYS1. In all cases, the forced PYRO-O plan performed significantly better than the other plans. The main reason for the improvement was the use of a partial sort of *lineitem* index entries as against a full sort. The final sort on *partkey* was not very expensive as only a few tuples needed to be sorted.

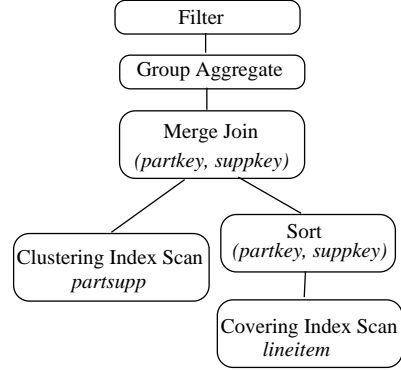
For Query 3 the plan generation phase (phase-1) was sufficient to select the sort orders and phase-2 does not make any changes. We shall now see a case for which phase-1 cannot make a good choice and the sort orders get refined by phase-2.

Experiment B2: This experiment uses Query 4, shown below, which has two full outer joins with two common attributes between the joins. We designed this experiment to see whether the systems we compare with exploit attributes common to multiple sort-based operators.

Query 4 *Attributes common to multiple joins*



(a) Default plan on SYS1



(b) Forced merge-join plan on SYS1 and default plan on SYS2

Figure 11: Plans for Query 3 (SYS1 and SYS2)




Figure 12: PostgreSQL




Figure 13: SYS1

```

SELECT * FROM R1 FULL OUTER JOIN R2
ON (R1.c5=R2.c5 AND R1.c4=R2.c4 AND R1.c3=R2.c3)
FULL OUTER JOIN R3
ON (R3.c1=R1.c1 AND R3.c4=R1.c4 AND R3.c5=R1.c5);

```

The tables R1, R2 and R3 were identical and each populated with 100,000 records. No indexes were built. As shown in Figure 14(a), both SYS1 and Postgres chose sort orders that do not share any common prefix. The plan chosen by PYRO-O is shown in Figure 14(b). In the plan chosen by PYRO-O, the two joins share a common prefix of $(c4, c5)$ and thus the sorting effort is expected to be significantly less. SYS2, not having an implementation of full outer join, chose a union of two left outer joins. The two left outer joins used to get a full outer join used different sort orders making the union expensive, illustrating a need for coordinated choice of sort orders.

Experiment B3: In this experiment we compare our approach of choosing interesting orders, PYRO-O, with the exhaustive approach, and a heuristic used by PostgreSQL. Postgres uses the

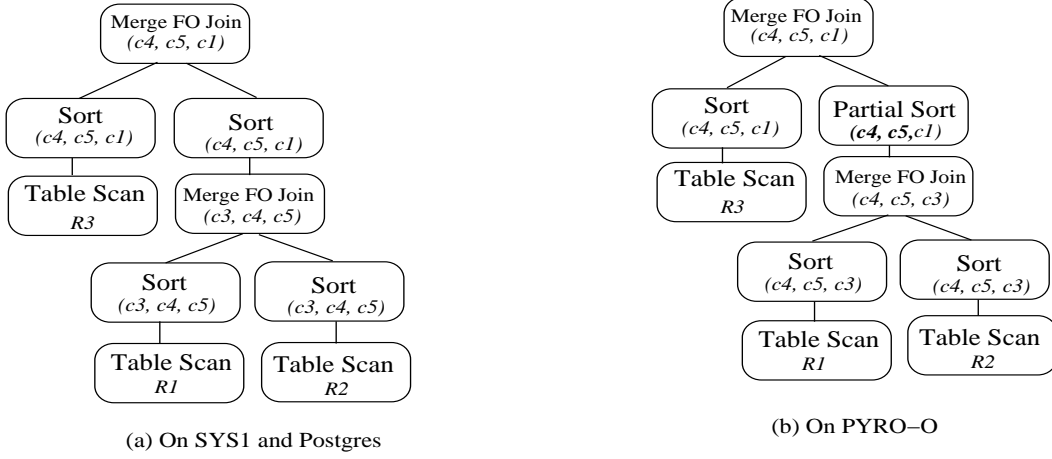


Figure 14: Plans for Query 4

following heuristic: for each of the n attributes involved in the join condition, a sort order beginning with that attribute is chosen; in each order, the remaining $n - 1$ attributes are ordered arbitrarily. We implemented Postgres’ heuristic in PYRO along with the extensions to exploit partial sort orders and call it PYRO-P. The exhaustive approach, called PYRO-E, enumerates all $n!$ permutations and considers partial sort orders. In addition, we also compare with PYRO, which chooses an arbitrary sort order, and a variation of PYRO-O, called PYRO-O⁻ that considers only exact favorable orders (no partial sort). Figure 15 shows the estimated plan costs. Note the logscale for y-axis. The plan costs are normalized taking the plan cost with exhaustive approach to be 100. In the figure, Q3 and Q4 are Query 3 and Query 4 of Experiments B1 and B2. Q5 and Q6 are Queries 5 and 6 shown below. For Q3 and Q4, as very few attributes were involved in the join condition, Postgres’ heuristic along with extensions to exploit partial sort orders, produced plans which were close to optimal. However, for more complex queries the heuristic does not perform as well since it makes an arbitrary choice for secondary orders.

Query 5 *Total value executed for a given order*

```
SELECT T1.UserId, T1.BasketId, T1.ParentOrderId, T1.WaveId, T1.ChildOrderId,
       (T1.Quantity * T1.Price) as OrderValue, SUM(T2.Quantity * T2.Price) as ExecutedValue
FROM TRAN T1, TRAN T2
WHERE T1.UserId=T2.UserId AND T1.ParentOrderId=T2.ParentOrderId AND T1.BasketId=T2.BasketId
      AND T1.WaveId=T2.WaveId AND T1.ChildOrderId=T2.ChildOrderId AND T1.TranType='New'
      AND T2.TranType='Executed'
GROUP BY T1.UserId, T1.BasketId, T1.ParentOrderId, T1.WaveId, T1.ChildOrderId;
```

Query 6 *Basket Analytics*

```
SELECT * FROM BASKET B, ANALYTICS A
WHERE B.ProdType = A.ProdType AND B.Symbol = A.Symbol AND B.Exchange = A.Exchange;
```




Figure 15: More Queries




Figure 16: Scalability

6.3 Optimization Overheads

The optimization overheads due to the proposed extensions were negligible. During plan generation, the number of interesting orders we try at each join or aggregate node is of the order of the number of indices that are useful for answering the query. We found that in most real-life queries this number is fairly small. Figure 16 shows the scalability of the three heuristics. For the experiment a query that joined two relations on varying number of attributes was used. Though PYRO-P and PYRO-O take the same amount of time in this experiment, in most cases, the number of favorable orders is much less than the total number of attributes involved and hence PYRO-O generates significantly fewer interesting orders than PYRO-P.

The plan-refinement algorithm presented in Section 4.2 was tested with trees up to 31 nodes (joins) and 10 attributes per node. The time taken was negligible in each case. The execution of plan refinement phase took less than 6 ms even for the tree with 31 nodes.

Both the optimizer extensions and the extension to external-sorting (MRS) were fairly straight forward to implement. The optimizer extensions neatly integrated into our existing Volcano style optimizer.

7 Conclusion and Future Work

We addressed the issue of choosing interesting sort orders. We showed that even a simplified version of the problem is *NP-hard* and proposed principled heuristics for choosing interesting orders. Our techniques take into account important issues such as partially matching sort orders and operators that require matching sort orders from multiple inputs. We presented detailed experimental results to demonstrate the benefits due to our techniques.

Secondary indices that do not cover the query can also be efficient to obtain a desired sort order if the actual tuple fetch can be deferred. Deferring the fetch until a point where the extra attributes are actually needed can be very effective when a highly selective filter discards many rows before the fetch is needed. Such optimizations must be made in a cost-based manner.

Next, we would like to investigate additional ways of speeding up nested iteration involving

user-defined functions. Though decorrelation is not always applicable for user-defined functions, we expect many cases that can be unnested using a combination of dataflow analysis and known unnesting methods.

In many cases users do not retrieve the whole of the query result but need only the first few tuples. In Top-K queries this interest is explicitly stated and the query optimizers attempt to find a plan that is optimal for returning the first K results. However, there are cases where the value of K is not known a priori. The user may decide to terminate the query execution at an arbitrary point, perhaps based on the results seen so far. We would like to study plans that adapt with the number of rows retrieved.

References

- [DPS02] Josep Diaz, Jordi Petit, and Maria Serna. A Survey of Graph Layout Problems. *ACM Comput. Surv.*, 34(3), 2002.
- [ECW92] Vladimir Estivill-Castro and Derick Wood. A survey of adaptive sorting algorithms. *ACM Comput. Surv.*, 24(4), 1992.
- [GM93] Goetz Graefe and W.J. McKenna. The Volcano Optimizer Generator: Extensibility and Efficient Search. In *ICDE*, 1993.
- [Gra93] Goetz Graefe. Query Evaluation Techniques for Large Databases. *ACM Comput. Surv.*, 25(2), 1993.
- [GRS05] Ravindra Guravannavar, H. S. Ramanujam, and S. Sudarshan. Optimizing Nested Queries with Parameter Sort Orders. In *VLDB*, pages 481–492, 2005.
- [Knu73] D. Knuth. *The Art of Programming, Vol. 3 (Sorting and Searching)*. Addison-Wesley, 1973.
- [Lar03] Per-Åke Larson. External sorting: Run formation revisited. *IEEE Trans. Knowl. Data Eng.*, 15(4), 2003.
- [NM04a] Thomas Neumann and Guido Moerkotte. A Combined Framework for Grouping and Order Optimization. In *VLDB*, 2004.
- [NM04b] Thomas Neumann and Guido Moerkotte. An Efficient Framework for Order Optimization. In *ICDE*, 2004.
- [SMD⁺79] P. Griffiths Selinger, M.M.Astrahan, D.D.Chamberlin, R.A.Lorie, and T.G.Price. Access Path Selection in a Relational Database Management System. In *Proceedings of ACM SIGMOD Conference*, 1979.
- [SSM96] David Simmen, Eugene Shekita, and Timothy Malkemus. Fundamental Techniques for Order Optimization. In *Proceedings of ACM SIGMOD Conference*, 1996.
- [WC03] Xiaoyu Wang and Mitch Cherniack. Avoiding Sorting and Grouping In Processing Queries. In *VLDB*, 2003.

A Optimality with Exact Favorable Orderes

Theorem A.1 *The set $I(e, o)$ computed with exact ford-mins contains an optimal sort order o_p for the optimize goal $e = (e_l \bowtie e_r)$ with (o) as the required output sort order.*

Assumption:

If o_1, o_2 are two orders such that $\text{attrs}(o_1) = \text{attrs}(o_2)$, $\text{cpu-cost}(e, o_1) = \text{cpu-cost}(e, o_2)$

Proof: We define the plan cost of an order o' for the optimization goal $(e = e_l \bowtie e_r, o)$, where o is the required output order as follows:

$PC(e, o, o') = \text{cbp}(e_l, o') + \text{cbp}(e_r, o') + \text{coe}(e, o', o) + CM$, where CM is the cost of the merge join operator, which we assume to be same for all o' such that $\text{attrs}(o') = S$, where S is the set of attributes involved in the join predicate.

Let o_b be an *optimal order* not belonging to set $I(e)$. We show that $\exists o_p \in I(e)$ such that $PC(o_p) = PC(o_b)$

Case 1: $o_b \notin \text{ford}(e_l) \cup \text{ford}(e_r)$

Let o_p be an order $\in I(e)$ such that $o \wedge S \leq o_p$. Note that there must exist such an order in $I(e)$ since $o \wedge S \in \mathcal{T}$ (Step 1 in computing $I(e)$)

$$PC(o_b) = \text{cbp}(e_l, o_b) + \text{cbp}(e_r, o_b) + \text{coe}(e, o_b, o) + CM$$

Since, $o_b \notin \text{ford}(e_l) \cup \text{ford}(e_r)$ we can write

$$PC(o_b) = \text{cbp}(e_l, \epsilon) + \text{coe}(e_l, \epsilon, o_b) + \text{cbp}(e_r, \epsilon) + \text{coe}(e_r, \epsilon, o_b) + \text{coe}(e, o_b, o) + CM$$

Since $\text{attrs}(o_b) = \text{attrs}(o_p)$, $\text{coe}(e_l, \epsilon, o_b) = \text{coe}(e_l, \epsilon, o_p)$. Hence,

$$\begin{aligned} PC(o_b) &= \text{cbp}(e_l, \epsilon) + \text{coe}(e_l, \epsilon, o_p) + \text{cbp}(e_r, \epsilon) + \text{coe}(e_r, \epsilon, o_p) + \text{coe}(e, o_b, o) + CM \\ &\geq \text{cbp}(e_l, o_p) + \text{cbp}(e_r, o_p) + \text{coe}(e, o_b, o) + CM \end{aligned}$$

Further, $o_b \wedge o \leq o_p \wedge o$ since $o \wedge S \leq o_p$ and o_b is a permutation of S . Hence,

$$\begin{aligned} PC(o_b) &\geq \text{cbp}(e_l, o_p) + \text{cbp}(e_r, o_p) + \text{coe}(e, o_p, o) + CM \\ PC(o_b) &\geq PC(o_p). \end{aligned}$$

Case 2: $o_b \in \text{ford}(e_l) \cup \text{ford}(e_r)$

Case 2A: $o_b \in \text{ford}(e_l)$ OR $o_b \in \text{ford}(e_r)$ but not both. Without loss of generality assume $o_b \in \text{ford}(e_l)$. This implies one of the following:

1. $\exists o' \in \text{ford-min}(e_l)$ such that $o_b \leq o'$

$$o' \in \text{ford-min}(e_l) \implies \exists o_p \in I(e) \text{ such that } o' \wedge S \leq o_p$$

Since o_b is a permutation of S we have $o_b \wedge S = o_b$

Therefore $o' \wedge S = o_b$ and hence $o_b \leq o_p$

Further, since $|o_b| = |o_p|$, we conclude $o_b = o_p$. This contradicts our earlier assumption that $o_b \notin I(e)$

2. $\exists o' \in \text{ford-min}(e_l)$ such that $o' \leq o_b$ and $\text{cbp}(e_l, o') + \text{coe}(e_l, o', o_b) = \text{cbp}(e_l, o_b)$
 $o' \in \text{ford-min}(e_l) \implies \exists o_p \in I(e)$ such that $o' \wedge S \leq o_p$

$$\begin{aligned} PC(e, o_b) &= \text{cbp}(e_l, o_b) + \text{cbp}(e_r, o_b) + \text{coe}(e, o_b, o) + CM \\ &= \text{cbp}(e_l, o') + \text{coe}(e_l, o', o_b) + \text{cbp}(e_r, o_p) + \\ &\quad \text{coe}(e, o_b, o) + CM \end{aligned}$$

(Note: $\text{cbp}(e_r, o_b) = \text{cbp}(e_r, o_p)$ as $o_b \notin \text{ford}(e_r)$)

Since $\text{attrs}(o_b) = \text{attrs}(o_p)$ we can write

$$\begin{aligned} PC(e, o_b) &= \text{cbp}(e_l, o') + \text{coe}(e_l, o' o_p) + \text{cbp}(e_r, o_p) + \\ &\quad \text{coe}(e, o_b, o) + CM \\ &\geq \text{cbp}(e_l, o_p) + \text{cbp}(e_r, o_p) + \text{coe}(e, o_b, o) + CM \end{aligned}$$

case (a) $o' \leq o$

Then we could choose o_p from $I(e)$ such that

$o \leq o_p$. Hence, $\text{coe}(e, o_b, o) \geq \text{coe}(e, o_p, o)$

$$\begin{aligned} PC(e, o_b) &\geq \text{cbp}(e_l, o_p) + \text{cbp}(e_r, o_p) + \text{coe}(e, o_p, o) + CM \\ &\geq PC(e, o_p) \end{aligned}$$

case (b) $o' \not\leq o$

Now, $o' \wedge o = o_b \wedge o = o_p \wedge o$ ($\because o' \leq o_b$ and $o' \leq o_p$)

$$\begin{aligned} PC(e, o_b) &\geq \text{cbp}(e_l, o_p) + \text{cbp}(e_r, o_p) + \text{coe}(e, o_p, o) + CM \\ &\geq PC(e, o_p) \end{aligned}$$

Case 2B: $o_b \in \text{ford}(e_l)$ as well as $o_b \in \text{ford}(e_r)$

This implies one of the following:

1. $\exists o' \in \text{ford-min}(e_l)$ (or $\text{ford-min}(e_r)$) such that $o_b \leq o'$
 In this case the proof can proceed as in **Case 2A(1)**.
2. $\exists o_1, o_2 \in \text{ford-min}(e_l)$ such that $o_i \leq o_b$ and $\text{cbp}(e_l, o_i) + \text{coe}(e_l, o_i, o_b) = \text{cbp}(e_l, o_b)$ (for both $i = 1$ and $i = 2$)
 Since $o_1 \leq o_b$ and $o_2 \leq o_b$, either $o_1 \leq o_2$ or $o_2 \leq o_1$. Hence, $\exists o_p \in I(e)$ such that $o_1 \wedge S \leq o_p$ and $o_2 \wedge S \leq o_p$
 Choosing this o_p and proceeding as in **Case 2A (2)** we can prove $PC(e, o_b) \geq PC(e, o_p)$

□